

UNITED STATES PATENT APPLICATION

FOR

COPPER-FACED MODULES, IMPRINTED COPPER CIRCUITS, AND THEIR
APPLICATION TO SUPERCOMPUTERS

INVENTORS:

PETER C. SALMON

BY

ALDO J. TEST
REGISTRATION NO. 18,048
DORSEY & WHITNEY LLP
4 EMBARCADERO CENTER, SUITE 3400
SAN FRANCISCO, CA 94111-4187
TELEPHONE: (650) 494-8700
FACSIMILE: (650) 494-8771

ATTORNEY DOCKET NO. 34091/AJT

EXPRESS MAIL NO. EV 182098341 US

RELATED APPLICATIONS

This application claims priority to provisional Application Serial No. 60/496,948 filed August 20, 2003.

BRIEF DESCRIPTION OF THE INVENTION

5 This invention relates to apparatus and method for building microelectronic modules having one or two conductive faceplates, imprinting methods, and their application to blade servers and supercomputers.

BACKGROUND OF THE INVENTION

10 Conventional printed circuit boards are constructed using copper signal planes laminated between glass-epoxy layers. Copper planes (foils) are also used for ground and for power supply voltages. For conventional epoxy laminate boards, trace and space dimensions are typically 100 microns each for a trace pitch of 200 microns, and hole diameters for plated through holes are typically 340 microns or more. More advanced boards are available from companies such as Unitive, Inc. of Research Triangle Park, North Carolina, USA. Unitive's boards employ BCB, a
15 spin-on resin available from Dow Chemical Company, rather than glass-epoxy. For a copper trace thickness of 2-3 microns, these boards achieve trace widths as small as 12 microns and trace spacing as small as 13 microns; traces on different levels are connected using vias with a diameter of 35 microns and a pitch of 45-55 microns.

20 At 10GHz the dielectric constant of BCB is 2.65 and its dissipation factor is 0.002. Its coefficient of thermal expansion (CTE) is 52 ppm/°C, and its moisture uptake is 0.14% by weight (as reported by Dow Chemical). Moisture uptake is critical because small amounts of absorbed water can substantially raise the effective dielectric constant, and a higher dielectric constant may effectively prohibit high frequency operation. The CTE of copper is 17 ppm/°C; various researchers list the CTE of silicon as 2.6-4.2 ppm/°C. The coefficients of thermal expansion are
25 important because the manufacturing process typically requires thermal cycles, and mismatches in CTE between materials in a stacked configuration result in mechanical stresses that can damage the board, or components attached to it. Additional thermal cycling is typically present during operation of the board, but usually this variation is less than during manufacturing and assembly.

In recent years, imprinting methods have been developed for several products. An example is the compact disc (CD) that has been manufactured by pressing a master tool into a plastic material, leaving an imprint of the desired pattern. A feature size of one micron and a fabrication cost less than a dollar per square foot have been reported. For both integrated circuit manufacture and printed circuit board manufacture, the imprinting method has potential to be an inexpensive patterning method compared with current photolithographic patterning methods. Photolithography has been the mainstay of integrated circuit patterning for many decades; it generally requires highly sophisticated tools for projecting a beam of light through a mask onto photosensitive materials. A fabrication facility to process silicon wafers by this method typically costs more than US\$1B today. By contrast, imprinting requires unsophisticated tools, yet it can produce fine features, smaller than 100 nanometers in some applications. The tools required for imprinting generally include a laminating press with provisions for aligning the layers, and for heating the thermoplastic material to be patterned. Cold embossing has also been developed using glass tools with embossed features etched therein, and UV curable dielectrics.

Electroplating is typically used to build up the copper layers, and chemical/mechanical polishing (CMP) is preferably used for planarizing each layer after plating. Imprinting and CMP methods can be combined to implement a dual damascene process using copper as the conductor material. Dual damascene processes are known in the art; “dual” refers to the fact that two depths of copper are implemented: trenches used for creating traces have a lesser depth than cylindrical holes for vias. The combination of tools for imprinting, plating and CMP may cost less than 1% of the cost of equivalent photolithographic tools. Other conductors may be used instead of copper, but copper offers a compelling combination of good electrical and thermal conductivity, adequate mechanical properties (especially in the form of dispersion strengthened copper, to be further discussed), and an infrastructure of existing tools and processes for drilling, electroplating, etching, and polishing at reasonable cost.

Liquid crystal display (LCD) panel fabrication plants are now being built for glass panels measuring 1870 x 2200mm. The thickness of these glass panels is 0.7mm, similar to the preferred thickness of 0.6mm for copper substrates of the current invention. This means that an infrastructure of semiconductor processing equipment, particularly including thin film coating and etching equipment could be adapted to handle copper substrates in panel sizes up to around 2 meters square. This may be useful for coating thin film adhesion layers and seed layers on large

panels of the current invention, and for plasma etching dielectric materials, as will be further described. The photolithographic patterning capability for large panels, generally employing step and repeat exposure systems ("steppers"), may also be used for fabricating the embossing tools described herein.

5 The preferred method of imprinting discussed herein uses an embossing tool in the form of a flexible foil, hereinafter called a "toolfoil". It is also possible to use rigid embossing tools, particularly if release agents are applied to the tool to aid in separation; also if the impressions are shallow they require a relatively small force for release. A release agent comprised of low surface energy material such as teflon is also preferably provided on the toolfoils discussed
10 herein. One method of making a toolfoil is to electroplate nickel in an additive process to create a master or "father" foil. The sidewalls of the plated features preferably have an angle of about 5° to the vertical. This release angle is useful so that negatives of fathers can be made to produce "mothers", and negatives of mothers are made to produce "sons", which are the equivalent of photo-tool working plates. Suitable toolfoils can be obtained from Tecan Components Ltd.,
15 Dorset, England. If the toolfoil is used to implement a dual damascene process, then two nickel thicknesses are required and two photo-tools (glass masks) will be used in the fabrication of the master.

 An alternative embodiment of the imprinting method may be used for large patterns that require a step and repeat methodology. Molecular Imprints of Austin, Texas, USA, has
20 developed a step and flash imprint lithography process called S-Fil. In this process, a substrate is coated with an organic planarization layer. Then a low viscosity photo-polymerizable imprint solution is dispensed on the surface. A surface treated transparent template bearing patterned relief structures is aligned over the coated substrate. The template is lowered into contact with the substrate, thereby displacing the solution, filling the imprint field, and trapping the photo-
25 polymerizable imprint solution in the template relief. Irradiation with UV light through the backside of the template cures the solution. The template is then separated from the substrate leaving a relief image on the surface that is a replica of the template pattern. A short halogen etch is used to clear any remaining thin webs of undisplaced material. A subsequent reactive ion etch into the planarization layer may be used to amplify the aspect ratio of the relief image.

30 Liquid crystal polymer (LCP) is a new dielectric material that has recently become available for imprinting applications in the printed circuit board arena. An example of this

material is R/Flex 3800 available from Rogers Corp., Circuit Materials Division, Chandler Arizona. It is available with a CTE matched to copper at 17 ppm/°C. Melting points of 280°C and 315°C are available, with thickness varying from 25 microns to 100 microns. From 1-10GHz the dielectric constant is 2.9 and the dissipation factor is 0.002. The moisture uptake is 0.04% by weight (compared with 0.14% for BCB), resulting in good stability for high frequency applications.

Dispersion strengthening is a method for improving the strength properties of copper, without seriously affecting its electrical and thermal conductivity. Cold rolled sheets of dispersion strengthened copper (DSC) known as Glidcop are available from SCM Metal Products, Inc., North Carolina. For use as a substrate for a printed circuit of the current invention, this material is available in thickness ranging from 125-625 microns. By incorporating minute amounts of aluminum oxide to pin the grain boundaries of the copper, the yield strength of DSC is typically improved by about 10 times, while the thermal and electrical properties are degraded by less than 1%.

Electroplating methods are well known in the art. Current processes support fabrication of via structures with aspect ratios (depth:diameter) as great as 10. Using layered plating solutions and sophisticated power supplies including reverse pulse biasing, void-free plated structures are achievable.

CMP is also well known in the art. A substrate to be polished is held in a polishing chuck so that typically one third of its edge dimension extends below the chuck. A polishing slurry is provided between the exposed surface and a rotating wheel having a finely textured surface; the substrate may simultaneously rotate and orbit in a planetary motion with respect to the wheel. The desired result is a polished planar surface with clearly defined copper features embedded in the dielectric resin.

Modern computer circuits such as multi-chip modules for computer server applications typically operate at GHz frequencies and with large power supply currents at low operating voltages: 200 amps at 1.0V is a typical requirement. Cooling of the module is a critical issue, and building such circuits on a copper substrate can help address the cooling requirements. In addition, a "copper sandwich" will be described having integral copper plates at both the top and bottom of the assembly for improved ruggedness and better thermal access to the heat-producing chips.

Flip chip assembly is generally recognized as the most advanced assembly method in terms of system density and performance. It enables bare integrated circuit (IC) chips to be assembled, in preference to packaged parts. The chips can have area arrays of input/output (I/O) bonding pads, rather than just at the chip periphery. Inductance of these chip-to-board connections is substantially lower than that of wire bonds, and power pins can be located close to the circuit blocks that need the power. Advanced flip chip assembly methods have recently been reported. One such method is to provide gold stud bumps on the IC chips, and corresponding wells filled with solder on the board. This structure supports pad pitches of 100 microns or less and also routine replacement of defective die using a rework process.

There have been two major impediments to the integration of large systems that are exclusively or primarily assembled using flip chip assembly methods: the inability to effectively test such a system (particularly a functional test at full system speed), and the inability to rework defective chips in the assembly. For these reasons, many flip chip assemblies have been limited to 10 IC chips or fewer, because the cost of scrapping defective assemblies becomes prohibitive with a greater number of chips. Solutions to these problems have been recently proposed. Firstly, a special-purpose test chip or chips may be provided on the board under test; working together with a test support computer this chip can provide the means to functionally test the module at full system speed. The test chip preferably includes high speed sampling circuits and comparators that are under control of the support computer. The support computer performs low speed testing chores such as boundary scan and loading of test files, and also hosts diagnostic software for aiding a test operator in determining which chips need to be replaced, if any. Secondly, the proposed variation of flip chip assembly allows effective rework of defective chips. In summary, the rework method is as follows. The board is placed on a hot plate and the temperature is raised to a level just below the melting point of the solder in the wells (the preferred melting point of the preferred In:Ag solder is 143°C). Then a rework wand emitting hot inert gas is directed at the backside of the defective chip; the solder in the wells melts for this chip, but not for neighboring chips that are not defective. Focused infra red systems have also been deployed for heating the area local to a single chip without melting the solder of surrounding chips. After the solder of the defective chip is molten, the stud bumps are withdrawn from the wells, the surface is inspected and cleaned as necessary, the wells are touched up with additional solder paste as required, and a replacement part is picked, flipped,

aligned, and inserted. After re-flowing the solder for the replacement part and validating the assembly with another module test, the rework cycle is complete. There is preferably no epoxy under layer beneath the defective chip (whose removal would be labor intensive, difficult, and potentially damaging to the board). Also, there are no delicate traces around or near the I/O pads that can be damaged during the rework process; the receiving terminal becomes the solder paste in the well rather than the pad. Finally, the materials used easily tolerate the rework temperatures. These factors result in a rework procedure that may be repeated as many times as necessary, enabling the integration of systems comprising hundreds or thousands of IC chips, assembled onto a single monolithic high performance substrate (or blade). This high level of integration in turn enables supercomputer architectures of the current invention.

The purpose of the epoxy under layer between chip and board is to prevent mechanical failure such as cracking that can arise from stresses accompanying temperature excursions that occur during manufacturing or operation. Part of the justification for eliminating this under layer in the proposed flip chip mounting structure is that the preferred arrangement of gold stud bumps inserted into wells filled with solder is mechanically stronger (shear forces can not easily detach a stud bump from its pad or a well from its pad). Because the stud bumps are formed from gold, and because gold is one of the most ductile materials, and because the proposed stud bumps have a pointed shape, the proposed structure is also more mechanically compliant than previous structures such as solder balls re-flowed onto matching lands. In addition, the low melting point of the proposed indium based solder (143°C) results in lower thermal strains than would occur with commonly used solders that melt at higher temperatures (63:37 Sn:Pb solder melts at 183°C). Even with this improved flip chip attachment it may be necessary to limit the maximum chip size in the proposed assembly structures, to limit the stress imposed.

For interconnecting modules, module access cables have been proposed that use a similar arrangement of stud bumps and wells as described for attaching the IC chips. These module access cables can support pin pitches of 100 microns or less, and should be re-workable using the same method as outlined for reworking the IC chips.

SUMMARY OF THE INVENTION

The current invention is intended to address the need for microelectronic assemblies that support operating frequencies of the order of 10GHz (data rates of the order of 10 Gbps), and

power supply currents of several hundred amps. An alternative application is to make lower performance electronic assemblies (sub 1GHz) less expensively than using current printed circuit board and assembly methods.

5 The preferred embodiment has a base copper layer for mechanical support. A method is described for fabricating feedthroughs in the copper substrate for signals and power. By alternately fabricating layers of dielectric resin and copper conductors, a printed circuit with multiple power and signal planes is built up. Imprinting is used to pattern the layers, preferably using nickel toolfoils. A special assembly layer is preferably fabricated on top of the interconnection circuit wherein a well filled with solder paste is provided at each I/O pad of the
10 board. IC chips are provided with a gold stud bump at each of their I/O pads, and the chips are assembled onto the board by inserting the stud bumps into the wells, then melting the solder to form mechanical and electrical connections. In the preferred embodiment, no epoxy under layer is used between the IC chips and the board. After all the chips are assembled, tested, and reworked as required, a top copper plate may be attached to the backsides of the assembled IC
15 chips, to make a robust mechanical, electrical and thermal package at the module level.

An alternate embodiment supports fabrication of modules such as multi-chip modules (MCMs) or system in package (SIP) devices that are integrated onto a motherboard. The motherboard may be a conventional board employing glass-epoxy laminate for example, or a large interconnection circuit fabricated on a copper substrate using methods described herein.
20 For either case, the copper substrate of the attached module is preferably provided with feeds for power and signals that pass through it and connect to the motherboard, typically on a 1 mm grid.

In the preferred embodiment, imprinting is used to pattern thermo-plastic dielectric layers, and copper conductors are employed. LCP is the preferred dielectric material and DSC is the preferred form of copper for the base layer. After CMP has been used to polish and planarize
25 the surface of a preceding layer, the following sequence summarizes the steps to form the next pair of interleaved dielectric and conducting layers:

- a) place a sheet of LCP on top of the polished planar assembly
- b) mount a toolfoil on a rigid carrier, align and position on top of the LCP sheet
- c) apply heat to soften the LCP
- 30 d) apply pressure to imprint the LCP
- e) cool to room temperature and separate toolfoil from its carrier

- f) peel the foil away from the assembly
- g) dry etch or sputter etch any remaining web of dielectric material, to expose the copper patterns underneath
- h) sputter deposit an adhesion layer such as Ti plus a seed layer of Cu
- 5 i) plate Cu to the desired thickness, typically 3-5 microns in the trenches, some of which is removed in the following step
- j) CMP to delineate the Cu patterns and planarize the surface

A variation of this imprinting procedure will be described for forming the special assembly layer having wells at each of the I/O pads of the interconnection circuit. "Imprinting" and "embossing" are used interchangeably in this application.

These techniques are employed to create copper-faced modules configured as blade servers, plus their integration into liquid-cooled supercomputers of the current invention.

BRIEF DESCRIPTION OF THE DRAWINGS

Preferred embodiments of the invention are described below with reference to the following accompanying drawings.

Figure 1 is a perspective view of a module built on a conductive substrate that uses a module access cable to communicate with other electronic systems.

Figure 2 is a perspective view of a module built on a conductive substrate that uses feeds through the copper plate (not shown) to communicate with a motherboard.

Figure 3 is a plan view of the underside of a conductive substrate having feedthroughs arranged in a two-dimensional grid.

Figure 4 is a perspective view of a module built on a conductive substrate that uses a wireless link to communicate with other electronic systems.

Figure 5 is a plan view of a system substrate having multiple modules attached, that connects to other electronic systems using a module access cable.

Figure 6 is a perspective view of a module having two conductive faces that connects to other electronic systems using a module access cable.

Figure 7 is a perspective view of a module having two conductive faces that connects to other electronic systems using feeds (not shown) through the bottom conductive face (substrate).

Figure 8 is a perspective view of a module having two conductive faces that communicates with other electronic systems by wireless means.

Figure 9 is a plan view of a system substrate having two conductive faces, with multiple modules attached to the bottom face (substrate), that connects to other electronic systems using a module access cable.

Figure 10 is a schematic side view of a stack of layers to be laminated for drilling.

Figure 11 is a cross-sectional view of a fragment of a conducting substrate having feedthrough holes drilled therein.

Figure 12(a) - (g) includes fragmentary cross-sectional views that illustrate the process steps required for fabrication of feedthroughs in a conducting substrate.

Figure 13A is a plan view of a fragment of an embossing tool in the form of a toolfoil showing a trench and an associated via.

Figure 13B is a cross-section of the trench and via of the toolfoil shown in Fig. 13A.

Figure 13C is a schematic cross-sectional view of an imprint corresponding to the toolfoil of Figures 13A and 13B.

Figure 14 is a schematic cross-sectional view of an aligner/laminating press for making imprinted patterns of the current invention.

Figure 15(a)-(d) illustrates in cross-section the process steps to fabricate a pair of interconnection layers including a conductive layer and a dielectric layer.

Figure 16(a) - (e) illustrates in cross-section the process steps to fabricate a special assembly layer including wells filled with solder.

Figure 17 is a cross-section of a fragment of a module assembly including a conductive substrate with feedthroughs, interconnection circuits, and a flip-chip mounted die.

Figure 18 is a flow chart that summarizes the process steps for imprinting a pair of layers of an interconnection circuit.

Figure 19 is a flow chart that summarizes the process steps for imprinting a special assembly layer, including wells for flip chip assembly.

Figure 20 is a schematic cross-sectional side view of a supercomputer of the current invention.

Figure 21 is a schematic top view of the supercomputer of Fig. 20.

Figure 22 shows a group of integrated circuit chips including processing, memory, and communication capabilities.

Figure 23 shows a supergroup of integrated circuit chips including multiple groups defined as in Fig. 22 plus other special purpose chips.

5 Figure 24 shows a blade layout including multiple super groups of integrated circuit chips, plus blade access ports.

Figure 25A shows a fragment of a blade computer with a blade access port.

Figure 25B shows cross-sectional details corresponding to section BB of Figure 25A.

Figure 26 shows a blade access cable attached to a rigid carrier using a release layer.

10 DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 1 shows an electronic module 10 having a conductive substrate 11. (base layer), and multiple integrated circuit chips (IC chips) such as 12 flip chip mounted thereon. Module 10 receives power and communicates with other electronic systems via module access cable 13, utilizing module access port 14, as will be further described. The attachment of IC chips 12 and
15 access cable 13 preferably employ a new version of flip chip assembly wherein each input/output pad of each chip and each cable preferably has a stud bump attached, and each stud bump mates with a well filled with solder formed on top of the interconnection circuit (not shown) fabricated on substrate 11. Detailed manufacturing steps for fabricating these circuits will be described. IC chip 15 may be a test chip, as will be further described. The material of conductive substrate 11
20 is preferably copper or a dispersion-strengthened copper (DSC). Figure 1 is an example of a module with a single conductive face.

Figure 2 shows another module 20 with a single conductive face 21 (substrate), except that this substrate has feeds for signals and power (feedthroughs) running through it (not shown), as will be further described. All of the conductive faces in modules and systems described in this
25 application, including substrates and top plates (to be described), are preferably manufactured from copper or DSC. Multiple IC chips such as 12 of Figure 1 are flip chip mounted as shown.

Figure 3 shows the backside of substrate 21 including a two-dimensional array of input/output terminals such as 22, preferably arranged on a 1mm grid. The fabrication of the feedthroughs and terminals will be further described.

Figure 4 shows another module 30 having a single conductive face 11 and multiple IC chips such as 12 flip chip mounted thereon as in Figure 1, and includes a wireless transceiver chip 32 (or group of chips), that preferably communicates via a two-way radio link 33 with other electronic systems. Power to module 30 may be provided using energy storage devices such as batteries or fuel cells (not shown), or via a cable like 13 of Figure 1, or via feedthroughs to a motherboard as in module 20 of Figure 2.

Figure 5 shows an electronic system 40 including a system substrate 41 (base plate) and multiple modules such as 20 of Figure 2 mounted thereon. Additional IC chips such as 43 may be included on substrate 41. Chips 43 may be used for testing, for power conversion, for maintenance and administration, for implementing interface protocols and driving communication interfaces, or for diagnostic purposes, as examples. System 40 interfaces to other electronic systems using system access cable 44 connected to system access port 45. However, system 40 may additionally or alternatively communicate with other systems using feeds through substrate 41 to another electronic assembly, or by wireless means.

Figure 6 shows electronic module 50 having two conductive faces: substrate 11 as in Figure 1 and top plate 52. Sandwiched between the two faces are flip chip mounted components such as 53 shown in dotted outline for illustration purposes. Module 50 may communicate with other electronic systems via a module access cable like 13 of Figure 1. After flip chip assembly (direct attachment) of all of the components onto substrate 11, a grinding or lapping step may be employed to make the components of uniform thickness, and provide a planar surface for attachment of top plate 52. Top plate 52 may be bonded or thermally coupled to the backsides of the IC chips using conductive epoxy or eutectic bonding materials or thermal grease, as is known in the art.

Figure 7 shows electronic module 60 having two conductive faces: a substrate 21 as in Figure 2 and a top plate 62. The backside of substrate 21 includes an array of feedthroughs for signals and power, as depicted in Figure 3.

Figure 8 shows module 70 having two conductive faces: substrate 11 as in Figure 1 and top plate 72. However in contrast with top plate 62 of Figure 7, top plate 72 has an opening 73 for radio waves to interact with radio transceiver chip 74, to enable a radio link 33 as in Figure 4.

Figure 9 shows electronic system 80 that is similar to system 40 except that a top plate covers the modules such as 20 of Figure 2 and IC chips such as 43 of Figure 5, shown in dotted

outline. System 80 communicates with other systems via a system access cable like 44 of Figure 5, but other communication methods may be additionally or alternatively employed.

Figure 10 shows a stack of metal layers 90 that can be assembled with layers of wax 91 between each of the metal sheets. This stack is used to drill feedthrough holes in the copper or DSC substrates that include feeds for signals and power. A suitable wax is Ocon-195, and a preferred wax thickness between the metal layers is around 0.008 inches or 200 microns. For the preferred substrate thickness of 600 microns, a suitable drill diameter is 0.0295 inches or 750 microns. A 0.040 inch thick layer of aluminum is preferably provided as a top layer 92 and as a bottom layer 93 of the stack; this provides more cleanly drilled holes at the inner layers, with some tearing permitted at the entry and exit, confined to the aluminum layers. In addition, the interposed wax layers 91 allow the drill to clear between each of the inner layers. A center drill is used to "spot drill" each location with a small diameter hole that is accurately placed, e.g., using a "0" center drill that will create a "dimple" with .030 inch diameter. Typically a drilling stack will include 8 layers of 600 micron thick substrate material 94. The preferred drilling machine is a vertical milling center such as the Bridgeport VMC 760. Using the same machine setup, it may also be desirable to mill the outline of the stack to a standard wafer size; the standard shape provides compatibility with mask aligners that may be employed for the lamination step, to be further described.

Figure 11 shows a fragment 95 of substrate 21 as in Figure 2 in which holes 96 have been drilled with a preferred pitch, P, 97 of 1mm or 1000 microns, and a preferred diameter, D, 98 of .0295 inches or 750 microns.

Figure 12 summarizes a preferred process sequence for fabricating feedthroughs in a copper substrate such as 21 of Figure 2, starting with a substrate layer that has been drilled as in Figure 11. Figure 12(a) shows that the drilled holes 96 have been filled with a dielectric material 100 such as a liquid crystal polymer. Since the available thickness of LCP may be limited to 100 microns, it may be necessary to stack and laminate 7 or 8 sheets to reliably fill drilled holes 96. After lamination, the top and bottom surfaces are preferably polished using CMP to achieve planarization. Figure 12(b) shows that dielectric material 100 has been concentrically drilled with a preferred drill size of .0225 inches or 572 microns. This step may be performed with individual substrates, or with a stack of substrates that has been pre-aligned and bonded with wax. For alignment purposes, 4 separate alignment holes are provided outside the two

dimensional array of feedthrough holes, one in each quadrant (not shown). Figure 12(c) shows that a layer of LCP 102 has been laminated onto the bottom of the drilled substrate, and a coating has been deposited that coats the bottom 103 and the walls 104 of drilled aperture 101. Coating 104 includes an adhesion layer such as titanium to a preferred thickness of 300-400 Angstroms or .03-.04 microns, plus a seed layer of copper to a preferred thickness of 500-800 Angstroms or .05-.08 microns. Figure 12(d) shows electroplated copper 105 filling the holes. Advanced plating methods including layered chemicals in the plating bath and pulse-reversing power supplies to effect plating from the bottom up are preferred. An uneven surface 106 results, and this is polished using CMP to achieve the planarized top surface 107 shown in Figure 12(e). Figure 12(f) shows that a masking layer 108 has been deposited and patterned to expose laminated dielectric material 102 at the base of each feedthrough, and this layer has been removed by dry etching through the metal mask to expose electroplated copper material 105. A preferred material for masking layer 108 is aluminum. Finally, Figure 12(g) shows the result of coating the exposed copper with an adhesion layer of titanium plus a seed layer of gold, electroplating gold to form gold contacts 109 on the bottom side of each feedthrough, and etching away masking layer 108. In the preferred embodiment, copper substrate 110 includes copper feedthroughs in a copper substrate that is 0.6mm thick, having a pitch P, 97, of 1mm, with a gold contact at the base of each feedthrough.

Figure 13A shows a fragment of a toolfoil 120 including a first raised portion 121 for imprinting a trench pattern, and a second raised portion 122 for imprinting a via. Figure 13B shows profile 123 of the toolfoil features in cross-section. Trench depth d1, 124, and via depth d2, 125, are shown. Release angle θ , 126, is preferably around 5° so that molded parts can be separated from toolfoil 120. It is also preferable to coat toolfoil 120 with a release material having a low surface energy such as teflon. Figure 13C shows the imprinted pattern achieved by impressing toolfoil 120 with profile 123 into a plastic material such as a heated liquid crystal polymer 127. Typically, the smaller depth d1, 124 is faithfully transferred, and d3, 128 equals d1. However, the larger depth d2, 125 is typically imperfectly transferred due to the spreading characteristics of the plastic material, and d4, 129 < d2. For the preferred geometries of the current invention, d1 equals approximately 4 microns, d2 equals approximately 20 microns, d3 equals approximately 4 microns, and d4 equals approximately 16 microns. After plating and polishing, these geometries result in a trace thickness of approximately 2 microns, and a total via

depth of approximately 18 microns, with approximately 5 microns of web material 130 removed by dry etching to expose the underlying contact metal (not shown), that is intended to connect with the via metal.

Figure 14 is a schematic of a laminating press/alignment fixture 134 suitable for making imprints of the current invention. This schematic represents a modified version of a commercial mask aligner, the MA600 available from Suss Microtec. Normally this tool operates with a glass mask (photo tool) held in a mask holder, and a semiconductor wafer held in a wafer chuck. It provides alignment optics for aligning wafer to mask, controlled UV radiation from above for exposing photo resist and other UV sensitive materials, and controlled pressure between mask and wafer up to about 15 pounds per square inch. The modification relates to the addition of infrared heating 143 directed toward the wafer chuck from below, as will be further explained. For the current application, the mask is replaced by an assembly including a glass plate 135 and toolfoil 120 of Figure 13B bonded to plate 135 using a UV release layer 136. The glass plate is captured in position by mask holder 137. From the bottom side, quartz chuck 138 holds copper substrate 139 (preferably machined in the form of a semiconductor wafer) which may have some interconnection circuit layers 140 already formed on it. A sheet of thermoplastic dielectric material 141 is positioned over interconnection circuits 140, but it is trimmed so as not to cover alignment holes in copper substrate 139. An operator employs alignment optics 142 to align copper substrate 139 with toolfoil 120. Toolfoil 120 has alignment features in the form of apertures that are sized slightly larger than the alignment holes drilled in copper substrate 139. The alignment operator achieves alignment by centering the concentric circles of the two sets of alignment features. Note that "copper substrate", "conductive substrate", and "DSC substrate" are interchangeable in this application. After substrate 139 and toolfoil 120 are properly aligned, heat is applied to copper substrate 139 using infrared radiating lamps positioned in a space underneath quartz chuck 138. The wavelength of IR radiation 143 is carefully selected so that the radiation will not be absorbed by quartz chuck 138, but will be absorbed by copper substrate 139. As substrate 139 heats up, it will emit radiation that can be detected by a sensor below quartz chuck 138 for tracking its temperature during heating and cooling cycles (at temperatures above a certain minimum, depending on the wavelength of the IR radiation). A wavelength of 2.5 microns is suitable in the preferred embodiment. Heating of substrate 139 is monitored until a desired softening point or melting point of the dielectric material is reached. In the preferred

embodiment, 280°C is the melting point of the LCP dielectric. The preferred embossing temperature is around 200°C so that the LCP remains contained in the embossed layers during each imprint cycle. The required pressure for imprinting will vary with the particular patterns imprinted, the materials used, and the embossing temperature. A process of trial and error is required to determine the optimal pressure for a particular setup; the goal is to provide clean trenches and vias and leave only a thin web of dielectric material to be removed. The thickness of the web for the preferred embodiment is approximately 5 microns (using 25 micron LCP sheets). After imprinting, IR radiation 143 is turned off and the parts are allowed to cool, while force 144 is applied. Then UV exposure 145 is turned on to radiate the UV release material, to effect separation between toolfoil 120 and glass plate 135 as the pressure is released. Substrate 139 with the new imprint pattern is removed from fixture 134, and toolfoil 120 is peeled away.

Figure 15(a) shows a fragment of a pre-existing interconnection circuit 150 that has been planarized and has an exposed contact region 151 on a dielectric material 152. Following an imprinting sequence as described in reference to Figure 14, trenches 153 and vias 154 are imprinted as shown in Figure 15(b), with a thin web of dielectric material 155 remaining. Web 155 is removed by dry etching or sputter etching, and the profile is as shown in Figure 15(c). An adhesion layer is coated onto the surface, such as 300-400 Angstroms of titanium, followed by a seed layer of copper having a preferred thickness of 500-800 Angstroms. The seed copper is then electroplated from the bottom up, and the top surface is polished to planarize it and to isolate the trench conductors 156 shown in Figure 15(d). The bottom of via 157 is in intimate contact with contact region 151 for a low resistance connection. Because copper and LCP both have CTEs of 17 ppm/°C, there should be minimal distortion or warping of the substrate during manufacture of these layer pairs, particularly if copper is also used as the base substrate layer.

Figure 16 illustrates a modified imprinting sequence for forming a special assembly layer including wells filled with solder. Figure 16(a) shows a pre-existing interconnection circuit 160 that has been polished and planarized, exposing polished trace 161. Figure 16(b) shows an imprint of a well feature 162, positioned above a conducting trace, fabricated using the procedure described in reference to Figure 14. Dielectric 163 is preferably BCB that has been applied using a spin-on method. Figure 16(c) shows that the remaining web of dielectric material 164 has been removed, and the top surface has been coated with an adhesion layer of titanium plus nickel to a thickness of approximately 50 micro inches or 1.3 microns. Nickel coating 165

provides a diffusion barrier between the copper traces of the underlying structure and the solder materials to be provided in the well. After polishing and planarization the nickel coating is removed except for coating in the well 166, as shown in Figure 16(d). Figure 16(e) shows that well 166 has been filled with solder paste 167 to form a well filled with solder 169, and a gold stud bump 168 has been inserted. After a reflow cycle to melt solder paste 167, a permanent mechanical and electrical connection is achieved. The preferred solder paste 167 is 3% silver and 97% indium with a melting point of 143°C. Preferred dimensions for a well include a diameter of 55 microns and a depth of 15 microns. By using imprinting to pattern them, economic fabrication is achievable for creating wells at a pitch of 100 microns or less. For logic elements requiring a large number of input/output connections, many thousands or even millions of wells can be filled with paste using one pass of a squeegee. Including the cost of the imprinted layer and a material cost of approximately 0.00008 cents per well for Indalloy 290 at \$3.48/gm leads to a manufacturing cost of less than 0.02 cents per well. Coupled with an estimated cost of 0.03 cents per gold stud bump, high density flip chip bonds may be achieved at 0.05 cents per connection. In addition, these bump/well connections are anticipated to be re-workable using simple and effective procedures, without risk of damaging the board to which the components are flip chip mounted.

Figure 17 shows a scaled fragment of a flip chip assembly 170 using the method of the current invention. Copper substrate 110 of Figure 12 has feeds for signals and power as previously described, including substrate 94 of Figure 10, dielectric material 100, and electroplated copper 105. Interconnection circuit 171 includes 4 signal layers 172 and 3 power planes 173, plus a special assembly layer 174 including wells filled with solder 169. Stud bumps 168 of Figure 16 are attached to I/O pads on IC chips like 15 of Figure 1 and inserted into wells 169 at a preferred pitch of 100 microns. Because of the good dielectric properties of LCP, and using known methods for creating differential signal pairs having controlled impedance, an assembly represented by Figure 17 can operate at high frequencies. With copper and LCP both having CTEs of around 17 ppm/°C and silicon having a CTE of around 3.5 ppm/°C, stress will be induced in assembly 170 as the components cool after the flip chip assembly procedure. The strains arising from these stresses have to be taken up by mechanical compliance of the bump/well structure, plus some compliance of the LCP dielectric. If a copper top plate like 52 of Figure 6 is employed, everything except the silicon chips will be thermally matched. Since the

copper base plate and top plate are typically thicker and stronger than the silicon chips, strain will occur mostly in the silicon chips. However, this strain is acceptable for most IC chips, and a module having matched copper faceplates will not warp significantly during the temperature excursions of typical operating environments. It should be mechanically robust, easily cooled, and electrically quiet (low electromagnetic radiation).

Figure 18 is a summary flow chart of the aforementioned process for creating a pair of interconnection layers, including a conducting layer and a dielectric layer of an interconnection circuit.

Figure 19 is a summary flow chart of the additional process steps required to form a special assembly layer including wells that can be filled with solder.

The copper-faced modules described herein can be integrated into servers and supercomputers to provide greater functional density and better cooling than have been available to date. Figure 20 is a schematic side view of a supercomputer 200 of the current invention, formed in the approximate shape of a cube. Cooling chambers 201 have a planar shape with a coolant fluid entering at an inlet port 202 and exiting at an outlet port 203. Manifolds (not shown) distribute the incoming fluid flows 204 and the outgoing fluid flows 205 through ports such as 202 and 203. A blade component (or "blade system") 206 is constructed on a conducting substrate 207, and is sandwiched between a pair of cooling chambers 201. Substrate 207 is like an enlarged version of substrate 11 of Figure 1. IC chips such as 208 are flip chip mounted to blade component 206, preferably using bumps 168 and wells 169 as previously described in reference to Figure 16. Because the flip chip assembly method provides a high functional density on blade 206, and because cooling chambers 201 can be constructed to have a thin profile, new levels of system density can be achieved. With higher system density, signal paths are shorter and operating speeds are higher. Key factors that enable such large flip chip assemblies include new methods of testing and rework that are described in the background section of this application.

Figure 21 shows a schematic top view of supercomputer 200 of Figure 20, showing outlet ports 203. Blade access cables 211 are provided for each blade to connect to each of its neighbors, as will be further described. System input/output cables 212 and 213 are also provided.

Figure 22 shows that IC chips may be formed into groups 220, each group preferably containing a computing chip 221, a bus interface chip or other communications chip 222, and one or more memory chips 223.

Figure 23 shows that groups 220 may be integrated into supergroups 230, comprising
5 multiple copies of group 220 plus a single copy of special purpose chips such as may be provided for testing 231, power conversion 232, cross bar switching 233 between pairs of groups or supergroups (for high speed group-to-group communications), and also chips for diagnostics, maintenance and administration 234. Another important function for a special purpose chip may be scheduling of events in preparation for connecting a pair of computing nodes, using a cross
10 bar switch for example. If a failure occurs in a blade system at a commercial business (“in the field”), any required rework will be expensive. Rather than rework any failing elements of such a blade system, it is preferable to provide an administration function that keeps track of defective elements or groups, and automatically switches them out of operation if any defects occur. To detect failures, periodic health checks of all hardware elements may be conducted in the
15 background, using test chips such as 231. The impact of a failed group will be small because it represents a small fraction of the total functionality, and maintenance costs will be minimized.

Figure 24 shows a top view of blade 206 of Figure 20 including multiple supergroups 230 as described in Figure 23 arrayed thereon plus an upper blade access port 241 and a lower blade access port 242. Combining the two access ports on each blade with upper and lower blade
20 access cables like 211 of Figure 21, it can be seen that each blade can be connected to each of its neighbors. Cable driver/receiver circuits may be implemented on transceiver chips such as 243.

As previously discussed, a semiconductor manufacturing infrastructure exists for flat panels as large as 2 meters on a side. In principle, large toolfoils could be fabricated for imprinting such a large panel in one step; the photo-imaging of such a large toolfoil area could
25 be accomplished using a stepper exposure system, employing photolithographic exposure systems already in service for LCD manufacture. If this is done, the preferred dielectric material (LCP) can be used and high frequency circuits can be produced in large panels using a single imprint cycle. Alternatively, imprinting of large panels can be achieved using a step and repeat imprinting process like the S-Fil process already described.

Figure 25A shows a blade access cable like 211 of Figure 21 connecting to upper blade
30 access port 241 of Figure 24. Figure 25B shows an expanded cross-sectional view of section BB

of Figure 25A. An interconnection circuit 251 preferably includes multiple signal and power layers plus special assembly layer 174 of Figure 17. Stud bumps 168 as shown in Figure 16 connect between input/output pads 252 on cable 211 and corresponding wells filled with solder connecting to traces or nodes 253 of interconnection circuit 251. The pitch P 254 of these
5 bump/well connections is preferably 100 microns. Cable 211 preferably includes a signal layer 255 and two ground layers 256. The arrangement of Figure 25B supports a connection density of 10,000 connections per square centimeter (at a pad pitch of 100 microns), while also supporting controlled impedances for high frequency operation.

For pad pitches as small as 100 microns, as shown in Figure 25B, dimensional stability
10 during assembly of the parts is essential. Cable 211 includes conductive and dielectric materials that flex; it is a flexible circuit and typically does not exhibit good dimensional stability. Therefore, cable 211 must be supported on a rigid carrier with good dimensional stability until all the connections have been made and tested. Carrier 260 in Figure 26 provides this capability; in the preferred embodiment it is a glass substrate on which cable 211 is fabricated. Release
15 layer 261 is preferably formed from a UV release material, and release is preferably achieved using intense UV exposure through carrier 260.